



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/703,909	11/01/2000	Todd P. Guay	1958.1031-000	3199
58403	7590	03/29/2006	EXAMINER	
BARRY W. CHAPIN, ESQ. CHAPLIN INTELLECTUAL PROPERTY LAW, LLC WESTBOROUGH OFFICE PARK 1700 WEST PARK DRIVE WESTBOROUGH, MA 01581			PHAM, KHANH B	
		ART UNIT		PAPER NUMBER
		2166		
DATE MAILED: 03/29/2006				

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	09/703,909	GUAY ET AL.	
	Examiner	Art Unit	
	Khanh B. Pham	2166	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133).

Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) Responsive to communication(s) filed on 17 January 2006.
- 2a) This action is **FINAL**. 2b) This action is non-final.
- 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) Claim(s) 1-18 and 20-136 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) Claim(s) _____ is/are allowed.
- 6) Claim(s) 1-18 and 20-136 is/are rejected.
- 7) Claim(s) _____ is/are objected to.
- 8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) The specification is objected to by the Examiner.
- 10) The drawing(s) filed on _____ is/are: a) accepted or b) objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

- 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) All b) Some * c) None of:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) <input type="checkbox"/> Notice of References Cited (PTO-892)	4) <input type="checkbox"/> Interview Summary (PTO-413)
2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948)	Paper No(s)/Mail Date: _____
3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) Paper No(s)/Mail Date: _____	5) <input type="checkbox"/> Notice of Informal Patent Application (PTO-152)
	6) <input type="checkbox"/> Other: _____

DETAILED ACTION

Claim Rejections - 35 USC § 102

1. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

2. **Claims 27-38, 44-46, 48-52 are rejected under 35 U.S.C. 102(e) as being anticipated by Adya et al. (US 6,266,658 B1), hereinafter “Adya”.**

As per claim 27, Adya teaches a system for evaluating a plurality of candidate index sets for a workload in a database system (See abstract) comprising:

- “a workload evaluator which evaluates each statement within the workload using collected database statistics” at Col. 6 line 55 to Col. 7 line 10;

(Adya teaches the module 230 that parses each individual query in the workload to identify candidate indexes by utilizing index usage information [collected database statistics] provided by database server 225.)

- “an index solution evaluation which, responsive to the workload evaluator, evaluates each index in a candidate index set with respect to the workload, the candidate index solution being one of the plurality of candidate index set, each candidate index set derived from an index superset formed by the union of a

current index set and a proposed index set" at Col. 8 lines 1-20 and Col. 9 lines 1-20;

(Adya teaches the list of potential indexes comprises "all existing indexes" and "propose indexes". The cost of each indexed is evaluated to eliminated indexes exceeds a predetermined fraction of the total workload cost.)

- "a solution/rollup evaluator which, responsive to the index solution evaluator, evaluates the candidate index solution" at Col. 10 lines 35-55 and Fig. 7, 716;

(Adya teaches the step of evaluate the resulting configuration R' to eliminate unused indexes form R' to produce configuration R'')

- "a solution refiner which, responsive to the solution/rollup evaluator, generates at least one new candidate index solution" at Col. 11 lines 13-30 and Fig. 7, 717;

(Adya teaches the step of pruning of indexes in R'' to eliminate indexes with low benefit to generate a new index solution (i.e., configuration F))

- "the at least one new candidates index solution being incorporated into the plurality of candidate index sets" at Col. 10 lines 44-54.

(Adya teaches a plurality of candidate index sets (i.e., R' , R'') are used. For example, the candidate index set R' is used to produce R'' , and R'' in turn is used to produce final configuration F).

As per claim 28, Adya teaches the system of claim 27 as discussed above.

Adya also teaches: "the solution refiner generates at least one new candidate index solution by eliminating at least one index within the candidate index solution that does not adhere to user-imposed constraints" at Col. 11 lines 13-30.

(Adya teaches the step of pruning of indexes in R'' to eliminate indexes with low benefit to generate a new index solution (i.e., configuration F.))

As per claim 29, Adya teaches the system of claim 28 as discussed above.

Adya also teaches: "the constraint is a user-defined constraint" at Col. 9 lines 10-15 and Col. 11 lines 15-25.

As per claim 30, Adya teaches the system of claim 28 as discussed above.

Adya also teaches: "the constraints is a memory usage constraint" at Col. 10 lines 33-35.

As per claim 31, Adya teaches the system of claim 27 as discussed above.

Adya also teaches: "the solution refiner generates at least one new candidate index solution by eliminating at least one index on a small table under evaluation, and wherein the at least one index does not enforce an integrity constraint" at Col. 2 lines 55-65.

As per claim 32, Adya teaches the system of claim 27 as discussed above.

Adya also teaches: "the workload evaluator evaluates an execution plan created by an optimizer, the execution plan comprising, for each statement of the workload, an execution plan which represents a series of steps for executing the statement; the

workload evaluator further generating and recording statistics based on the evaluation of the execution plan" at Col. 6 line 55 to Col. 7 line 10.

As per claim 33, Adya teaches the system of claim 32 as discussed above.

Adya also teaches: "each execution plan is created based on available access paths" at Col. 2 lines 30-40.

As per claim 34, Adya teaches the system of claim 32 as discussed above.

Adya also teaches: "each execution plan is created based on statistics for at least one schema object accessed by the statement" at Col. 7 lines 1-10 and Col. 1 lines 40-65.

As per claim 35, Adya teaches the system of claim 34 as discussed above.

Adya also teaches: "the at least one schema object is a table" at Col. 1 lines 40-65.

As per claim 36, Adya teaches the system of claim 34 as discussed above.

Adya also teaches: "the at least one schema object is an index" at Col. 1 lines 40-65.

As per claim 37, Adya teaches the system of claim 32 as discussed above.

Adya also teaches: "the workload evaluator, for a table accessed by a statement under evaluation, identifies at least one index which would be used to retrieve data from the table upon an execution of the statement" at Col. 6 lines 55-65.

As per claim 38, Adya teaches the system of claim 32 as discussed above.

Adya also teaches: "the workload evaluator determined a cost of the execution plan" at Col. 7 lines 12-13.

As per claim 44, Adya teaches the system of claim 32 as discussed above.

Adya also teaches: "the statistics include an index usage" at Col. 7 lines 5-10.

As per claim 45, Adya teaches the system of claim 32 as discussed above.

Adya also teaches: "the statistics include a cost of the execution plan" at Col. 7 lines 10-15.

As per claim 46, Adya teaches the system of claim 27 as discussed above.

Adya also teaches: "the statements are SQL statements" at Col. 8 lines 10-15.

As per claim 48, Adya teaches the system of claim 27 as discussed above.

Adya also teaches: "the solution refiner is responsive to a predetermined maximum number of allowed indexes" at Col. 10 lines 17-18.

As per claim 49, Adya teaches the system of claim 27 as discussed above.

Adya also teaches: "the solution refiner is responsive to available storage space" at Col. 10 lines 10-15.

As per claim 50, Adya teaches the system of claim 27 as discussed above.

Adya also teaches: "the proposed index set is provided by a user" at Col. 2 lines 40-45.

As per claim 51, Adya teaches the system of claim 27 as discussed above.

Adya also teaches: "the proposed index set is provided by an expert system" at Col. 2 lines 60-62.

As per claim 52, Adya teaches the system of claim 27 as discussed above.

Adya also teaches: "an execution plan is created without creating indexes which are not in the current index set" at Col. 11 lines 1-12.

Claim Rejections - 35 USC § 103

3. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

4. **Claims 1-6, 8-12, 16, 18, 20, 22-24, 26, 53-54, 56, 58-60, 62-64, 66-70, 74, 76-77, 79-81, 83-86, 88-90, 92-96, 100, 102-103, 105-107, 109-113, 115-117, 119-123, 127, 129-130, 132-134, 136 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chaudhuri et al. (US 5,960,423A), hereinafter "Chaudhuri (I)", and in view of Chaudhuri et al (US 6,223,171), hereinafter "Chaudhuri (II)".**

As per claims 1,53, 54, 110, Chaudhuri (I) teaches a method and system for evaluating a plurality of candidate index set for a workload of database statements in a database system (See abstract) comprising:

- "forming an index superset from a union of a current index set and a proposed index set" at col.7 lines 56-63 and Fig. 3, elements 232, 233, 234.

(Chaudhuri (I) teaches the storage of what-if indexes [proposed indexes that are absent from the database server] and existing indexes [current index set]. The storage of the indexes serves as an index superset).

- "deriving a candidate index set from the index superset, the derived candidate index set being incorporated into the plurality of candidate index sets" at Col. 2 lines 25-50 and col. 7 line 64 to col. 8 line 10.

(Chaudhuri (I) teaches that the superset of proposed index sets and existing index sets are considered in evaluating index configurations [candidate index sets]. Chaudhuri (I) also teaches indexes selected from plurality of index configurations are used to build the set of candidate indexes).

- "analyzing collected database statistics based on the derived candidate index set" at Col. 5 line 64 to Col. 6 line 2, Col. 9 lines 3-11, and Col. 10 lines 13-25.

(Chaudhuri (I) teaches that the statistics information included in the database schema are collected from database server 220 on startup. Statistical information for "what-if index" are generated by reading all n pages of the database tables". The costs [generated statistics] for each candidate index configuration are determined and stored.)

Chaudhuri (I) does not teach a user interface that presents the collected statistics. However, Chaudhuri (II) teaches a similar method includes the step of presents the statistics to users at Fig. 6-8 and Fig. 10-15. It would have been obvious to one ordinarily skilled in the art at the time of the invention to present these statistics to the administrator or user because of the advantages of allowing the user to analyze the data. By presenting the statistics, the user can gain a better understanding of the costs

and the advantages of particular indexes with each workload. Through the analysis, the user can determine the potential impact and improvements in performance of a candidate index set on a particular workload as taught at Chaudhuri (II) col. 1 line 64 to col. 2 line 13.

Claims 53-54 refers to "repeatedly derives a candidate index set" which can be found at Chaudhuri (I) col. 8 lines 46-56.

As per claims 2, 58, 84, and 111, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 1, 53, 54, 110 as discussed above. Chaudhuri (II) also teaches: "deriving current index statistics for the workload responsive to the current index set, the presented generated statistics comprising the generated current index statistics" at Col. 10 lines 13-21 and Fig. 15.

As per claims 3, 59, 85, and 112, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 1, 53, 54, 110 as discussed above. Chaudhuri (II) also teaches: "repeatedly deriving a candidate index set and analyzing collected statistics based on the proposed index set" at col. 23 lines 48-67.

As per claims 4, 60, 86, and 113, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 3, 59, 85, 112 as discussed above. Chaudhuri (II) also teaches: " terminating the repeated execution when at least one candidate index solution is found that adheres to user-imposed constraints and no further indexes can be removed from said candidate index solution without degrading performance of the workload and without degrading disabling an integrity constraint" at Col. 15 lines 35-64 and Col. 23 line 25 to Col. 24 line 6.

As per claims 6, 64, 90, and 117, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 1, 53, 54, 110 as discussed above. Chaudhuri (I) also teaches: "wherein analyzing statistics for a statement comprises generating at least one statistic based on an execution plan created by an optimizer" at col. 10 lines 53-56 and col. 8 lines 4-7.

As per claims 8, 66, 92, and 119, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 6, 64, 90, and 117 as discussed above. Chaudhuri (I) also teaches: "wherein the execution plan is based on statistics for at least one schema object accessed by the statement" at col. 10 lines 53-56 and col. 5 lines. 52-56. (The query optimizer may user an index over any single table of the database. The costs include the index configuration of both the indexes and the tables. The statistics are based upon the schema objects).

As per claims 9, 67, 93, and 120, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 8, 66, 92, and 119 as discussed above. Chaudhuri (I) also teaches: "the at least one schema object is a table" at col. 12 lines 20-29.

As per claims 10, 68, 94, and 121, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 8, 66, 92, and 119 as discussed above. Chaudhuri (I) also teaches: "the atleast one schema object is an index" at col. 10 lines 54-55.

As per claims 11, 69, 95, and 122, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 6, 64, 90, and 117 as discussed above. Chaudhuri (I) also teaches: "for a table accessed by a statement under evaluation, using the execution plan to identify at least one index that would be used to retrieve data from the table upon an execution of the statement" at col. 6 lines 9-13 and col. 10 lines 53-56. (The execution plan is returned with a cost estimate of executing a designated query for the designated candidate index configuration. The index selection tool attempts to select an index configuration that is optimal, so the cost will lead to the identification of an index).

As per claims 12, 70, 96, and 123, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 6, 64, 90, and 117 as discussed above. Chaudhuri (I) also teaches: "the optimizer generates a cost of the execution plan" at col. 10 lines 48-56. (The query optimizer returns an execution plan with a cost estimate. Because the optimizer returns both the execution plan and the cost estimate, the optimizer is considered to have generated both the plan and the cost).

As per claims 16, 74, 100, and 127, Chaudhuri (I), Chaudhuri (II) teach the method, computer program and system of claims 6, 64, 90, 117 as discussed above. Chaudhuri (II) also teaches: "the collected database statistics comprise the number of executions of the statement" at Figs. 6-7.

As per claims 18, 76, 102, and 129, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 6, 64, 90, and 117 as discussed above. Chaudhuri (II) also teaches: "the collected database statistics comprise an index usage" at Fig. 8.

As per claims 20, 77, 103, and 130, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 1, 53, 54, and 110 as discussed above. Chaudhuri (I) also teaches: "wherein the statements are SQL statements" at col. 5 lines 40-41 and col. 5 lines 61-64.

As per claims 22, 79, 105, and 132, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 1, 53, 54, and 110 as discussed above. Chaudhuri (I) also teaches: "deriving a candidate index set is responsive to a predetermined maximum number of allowed indexes" at col. 13 lines 1-9.

As per claims 23, 80, 106, and 133, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 1, 53, 54, and 110 as discussed above. Chaudhuri (II) also teaches: "wherein deriving a candidate index set is responsive to available storage space" at col. 1 lines 24-34.

As per claims 24, 81, 107, and 134, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 1, 53, 54, and 110 as discussed above. Chaudhuri (II) also teaches: "the proposed index set is provided by a user" at col. 12 lines 1-10.

As per claims 26, 83, 109, and 136, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 1, 53, 54, and 110 as discussed

above. Chaudhuri (I) also teaches: "an execution plan is created without creating indexes which are not in the current index set" at col. 7 lines 52-63.

As per claims 56, 62, 88, and 115, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 1, 53, 54, and 110 as discussed above. Chaudhuri (I) also teaches: "generating baseline statistics for each statement in the workload, wherein generating statistics is additionally based on the baseline statistics" at col. 6 lines 14-22 and col. 10 lines 48-56.

(Chaudhuri (1) teaches that the query optimizer estimates a cost for each query [statement] in the workload. The estimates for each query are then summed to determine statistics of the workload with the candidate index set).

As per claims 5, 63, 89, and 116, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 56, 62, 88 and 115 as discussed above. Chaudhuri (I) also teaches: "analyzing the collected baseline statistics comprises disabling current indexes" at Col. 14 lines 60-65.

5. Claims 7, 65, 91, and 118 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chaudhuri (I) and Chaudhuri (II) as applied to claims above, and further in view of Jakobsson et al. (US 5,924,088 A), hereinafter "Jakobsson".

As per claims 7, 65, 91, and 118, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 6, 64, 90, and 117 as discussed above. Chaudhuri (I) also teaches that an execution plan is created at col. 10 lines 53-56. Chaudhuri (I) also teaches that the indexes are used in order to access the database at col. 3 lines 50-54. Chaudhuri I and II do not explicitly teach that the "execution plan is based on available access paths". However, Jakobsson teaches a similar method for index selection includes the use of access paths at col. 4 lines 1-4. It would have been obvious to one ordinarily skilled in the art at the time of the invention to base the execution plan on an access path because the performance characteristics for a data retrieval may vary greatly depending on the choice of index access path as taught at col. 4 lines 8-12. The costs of the statements of the workload can be calculated through the use of the access paths. Using these paths would allow for improved calculations of the costs of each statement.

6. **Claims 13-15, 71-73, 97-99, and 124-126 are rejected under 35 U.S.C. 103(a)** as being unpatentable over Chaudhuri (I) and Chaudhuri (II) as applied to claims above, and further in view of Eberhard et al. (US 6,003,022), hereinafter "Eberhard".

As per claims 13, 71, 97, and 124, Chaudhuri (I) and Chaudhuri (II) teach the method, computer program and system of claims 12, 38, 70, 98, and 123 as discussed above. Chaudhuri (I) describes an invention that selects an index that would work best for a workload as described above. Chaudhuri (I) also teaches that each query does have a particular cost associated with each particular index at col. 10 lines 14 9.

Chaudhuri (I) and Chaudhuri (II) do not explicitly teach that the cost of the execution plan is derived from a resource as claimed. However, Eberhard does teach that the execution costs are derived from a resource at col. 3 lines 43-45. It would have been obvious to one ordinarily skilled in the art at the time of the invention to derive the cost of the execution plan from a particular resource because of execution costs in both CPU and I/O as taught at col. 3 lines 43-59. By calculating the costs of the queries under a particular index, the invention would be capable, of calculating how each candidate index can affect each statement in the workload. Deriving the cost of the execution plan from a resource use needed to execute the statement allows the system to know how each query will affect the system's performance. Considering the system's performance will allow a better cost to be derived for the indexing system.

As per claims 14, 72, 98, and 125, Chaudhuri (I), Chaudhuri (II), and Eberhard teaches the method, computer program and system of claims 13, 71, 97, 124 as discussed above. Eberhard also teaches: "the resource use includes CPU execution time" at col. 3 lines 43-44.

As per claims 15, 73, 99, and 126, Chaudhuri (I), Chaudhuri (II), and Eberhard teaches the method, computer program and system of claims 13, 71, 97, 124 as discussed above. Eberhard also teaches: "the resource use includes input/output access" at col. 3 lines 43-44.

7. Claims 39-41 are rejected under 35 U.S.C. 103(a) as being unpatentable over Adya as applied to claims 27-38, 44-46, 48-52 above, and further in view of Eberhard.

As per claim 39, Adya teaches the system of claim 38 as discussed above.

Adya does not explicitly teach that "the cost of the execution plan is derived from a resource" as claimed. However, Eberhard does teach that the execution costs are derived from a resource at col. 3 lines 43-45. It would have been obvious to one ordinarily skilled in the art at the time of the invention to derive the cost of the execution plan from a particular resource because of execution costs in both CPU and I/O as taught at col. 3 lines 43-59. By calculating the costs of the queries under a particular index, the invention would be capable, of calculating how each candidate index can affect each statement in the workload. Deriving the cost of the execution plan from a resource use needed to execute the statement allows the system to know how each query will affect the system's performance. Considering the system's performance will allow a better cost to be derived for the indexing system.

As per claim 40, Adya and Eberhard teach the system of claim 39 as discussed above. Eberhard also teaches: "the resource use includes CPU execution time" at col. 3 lines 43-44.

As per claim 41, Adya and Eberhard teach the system of claim 39 as discussed above. Eberhard also teaches: "the resource use includes input/output access" at col. 3 lines 43-44.

8. **Claims 17, 75, 101 and 128 are rejected under 35 U.S.C. 103(a)** as being unpatentable over Chaudhuri (I) and Chaudhuri (II) as applied to claims above, and further in view of Smith et al. (US 5,404,510 A), hereinafter "Smith".

As per claims 17, 75, 101, and 128, Chaudhuri (1) and Chaudhuri (II) teach the method, computer program and system include a query optimizer which searches for the best index to evaluate a given workload. Both Chaudhuri (I) and (II) do not explicitly teach that “the collected database statistics comprise a user-defined importance of the statement”. Smith does teach a user-defined importance of the statement at col. 7 lines 22-32. It would have been obvious to one ordinarily skilled in the art at the time of the invention to allow a user to define the importance of the statements found in the workload because the index selection tool could evaluate the statistics of the indexes according to the importance of the statements. A user could rank the importance of the queries in the workload. The ranking could then be used to define an index that works best for the workload. The index would then be created in consideration of the higher ranked statements.

9. **Claims 43 is rejected under 35 U.S.C. 103(a)** as being unpatentable over Adya as applied to claims 27-39, 45-46 and 48-52 above, and further in view of Smith et al. (US 5,404,510 A), hereinafter “Smith”.

As per claims 43, Adya teaches the system of claim 32 discussed above. Adya does not explicitly teach “the statistics include a user-defined important of the statement. Smith does teach a user-defined importance of the statement at col. 7 lines 22-32. It would have been obvious to one ordinarily skilled in the art at the time of the invention to allow a user to define the importance of the statements found in the workload because the index selection tool could evaluate the statistics of the indexes

according to the importance of the statements. A user could rank the importance of the queries in the workload. The ranking could then be used to define an index that works best for the workload. The index would then be created in consideration of the higher ranked statements.

10. **Claims 21, 78, 104, and 131 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chaudhuri (I) and Chaudhuri (II) as applied to claims 20, 77, 103 and 130 above, and further in view of Celis et al. (US 6,021,405), hereinafter "Celis".**

As per claims 21, 78, 104, and 131, Chaudhuri (I) teaches the use of statistics to find the cost of statements using a particular index. Both Chaudhuri I and II do not teach that the workload is reduced into unique statements in order to determine these costs. Celis teaches that the workload is reduced into unique expressions at col. 5 lines 37-44. Chaudhuri (I) teaches that the index selection tool is expensive and that each request to run query optimizer increases the cost of time and memory at col. 6 lines 23-35. Furthermore, Chaudhuri (I) teaches that the queries of a workload are sent to the query optimizer at col. 10 lines 13-60. Chaudhuri (I) also teaches that the cost evaluation tool attempts to reduce the number of invocations of query optimizer by determining costs of queries of workloads without invoking query optimizer at col. 10 lines 61-67. By not sending particular queries to the query optimizer, the workload is reduced. In particular, Chaudhuri (I) teaches the elimination of running repetitions of the query on atomic indexes. By not running repetitions of a query on similar indexes, Chaudhuri (I) is reducing the workload into unique statements. Celis teaches in greater

depth that redundant expressions are removed from the queries in order to save cost as taught at col. 1 lines 38-51.

It would have been obvious to one ordinarily skilled in the art at the time of the invention to generate the statistic using only unique expressions because of the ability to reduce processing time for redundant queries as taught at Celis col. 1 lines 41-51. Celis's invention is a query optimizer that is used to reduce the cost of each statement as taught at col. 1 lines 26-31. By only optimizing the unique statements, the optimizer will not be burdened by redundant statements. Chaudhuri (I) teaches the reduction of the number of invocations of query optimizer at col. 6 lines 51-59. By reducing the workload into unique statements, Chaudhuri (I) will invoke query optimizer fewer times for statements that had previously been optimized. The needless optimizations cost more processing time and ultimately increase the execution time for the query as taught at Cells col. 1 lines 42-51 and Chaudhuri I col. 6 lines 51-59.

11. **Claim 47 is rejected under 35 U.S.C. 103(a)** as being unpatentable over Adya as applied to claims 27-39, 45-46, 48-52 above, and in view of Celis.

As per claim 47, Adya teaches the system of claim 27 discussed above. Adya does not explicitly teach that "the workload is reduced into unique statements". Celis teaches that the workload is reduced into unique expressions at col. 5 lines 37-44. Celis teaches in greater depth that redundant expressions are removed from the queries in order to save cost as taught at col. 1 lines 38-51. It would have been obvious to one ordinarily skilled in the art at the time of the invention to generate the statistic using only

unique expressions because of the ability to reduce processing time for redundant queries as taught at Celis col. 1 lines 41-51. Celis's invention is a query optimizer that is used to reduce the cost of each statement as taught at col. 1 lines 26-31. By only optimizing the unique statements, the optimizer will not be burdened by redundant statements. By reducing the workload into unique statements, Adya will invoke query optimizer fewer times for statements that had previously been optimized. The needless optimizations cost more processing time and ultimately increase the execution time for the query as taught at Celis col. 1 lines 42-51.

12. Claims 25, 82, 108, 135 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chaudhuri (I) and Chaudhuri (II) as applied to claim 1 above, and further in view of Adya.

As per claims 25, 82, 108, and 135, Chaudhuri (I) and Chaudhuri (II) teaches the method, computer program and system as discussed in claims 1, 53, 54, 110 above. Both Chaudhuri (I) and Chaudhuri (II) do not specifically teach an expert system that chooses the index configuration. However, using an expert system to propose index set is well known in the art, as exemplary by Adya at Col. 2 lines 60-63. Thus, It would have been obvious to one ordinarily skilled in the art at the time of the invention to use an expert system because of the ability of expert systems to make index selections. By using an expert system, the administrator would not have to determine the most useful index set. The expert system could use prior knowledge and the calculations found in Chaudhuri (I) to determine an index set.

13. Claims 55, 61, 87 and 114 rejected under 35 U.S.C. 103(a) as being unpatentable over Chaudhuri (I) and Chaudhuri (II) as applied to claims 1, 53, 54, 110 above, and further in view of Gurry et al. (“Oracle Performance Tuning”), hereinafter “Gurry”.

As per claims 55, 61, 87, and 114, Chaudhuri (I) teaches the estimation of queries over candidate index sets as taught at col. 10 lines 48-56. Chaudhuri (I) and Chaudhuri (II) do not teach that the index volatility is included in the statistics. Gurry does teach the index volatility in the statistics at pages 353-354. Gurry teaches the volatility of indexes through the HEIGHT and DEL LF ROWS columns found in the statistics. If the statistics for the HEIGHT or DEL LF ROWS exceed a particular number, the index is a candidate for dropping and recreating. The dropping and recreating of an index is a part of the volatility of the index. It would have been obvious to one ordinarily skilled in the art at the time of the invention to include the index volatility in the statistics because of the costs incurred by dropping and recreating indexes. The index volatility demonstrates to the user the possibility of the index to be dropped and recreated. The volatile indexes may need to be rebuilt in order to obtain optimal performance [pg. 354]. Indexes over volatile tables may also become a space management problem [pg. 354]. The volatility statistics allow a DBA to be aware of the problems that may arise over particular index selections.

14. Claim 57 is rejected under 35 U.S.C. 103(a) as being unpatentable over Adya as applied to claims above, and further in view of Gurry.

As per claim 57, Adya teaches the system of claim 32 discussed above. Adya does not teach that the index volatility is included in the statistics. Gurry does teach the index volatility in the statistics at pages 353-354. Gurry teaches the volatility of indexes through the HEIGHT and DEL LF ROWS columns found in the statistics. If the statistics for the HEIGHT or DEL LF ROWS exceed a particular number, the index is a candidate for dropping and recreating. The dropping and recreating of an index is a part of the volatility of the index. It would have been obvious to one ordinarily skilled in the art at the time of the invention to include the index volatility in the statistics because of the costs incurred by dropping and recreating indexes. The index volatility demonstrates to the user the possibility of the index to be dropped and recreated. The volatile indexes may need to be rebuilt in order to obtain optimal performance [pg. 354]. Indexes over volatile tables may also become a space management problem [pg. 354]. The volatility statistics allow a DBA to be aware of the problems that may arise over particular index selections.

15. Claims 42 is rejected under 35 U.S.C. 103(a) as being unpatentable over Adya as applied to claims 27-38, 44-46, and 48-52 above, and in view of Finkelstein et al. (“Physical Database Design for relational Database), hereinafter “Finkelstein”.

As per claim 42, Adya teaches the system of claim 32 as discussed above. Adya does not explicitly teach that “the statistics include the number of execution of the statement. However, Finkelstein teaches a similar system for index selection utilizing collected statistics from database, wherein the statistics include “the number of

execution of the statement" at page 99. Thus, it would have been obvious to one of ordinary skill in the art at the time of the invention was made to modify Adya's teaching to include "the number of execution of the statement" in the statistics so that the workload could be calculated faster and more precisely because the same statement executed multiple times could be grouped and calculated at once.

Response to Arguments

16. Applicant's arguments filed January 17, 2006 have been fully considered but they are not persuasive. The examiner respectfully traverses applicant's arguments.

Regarding the 102 rejection based on Adya, applicant argued that Adya does not teach or suggest "candidate index set derived from an index superset formed by the union of a current index set and a proposed index set". On the contrary, Adya teaches a list of potential indexes, in which "certain restrictions are applied to reduce the number of potential indexes consider" (Col. 7 lines 39-40). Adya teaches the restrictions include: "3. All exiting indexes are retained" and "4. Indexes are proposes only for the most expensive queries..." (Col. 8 lines 4-6). Therefore, after applying the restriction, Adya's list of potential indexes, (i.e., "candidate index set") comprise: "all existing indexes" and "proposes indexes" as claimed.

Applicants argued that Adya's "potential index set" is different from applicant's "candidate index set" because Adya parses queries to identify candidate indexes to be placed in the set of candidate indexes, then applied the restriction to remove certain undesirable indexes. The examiner respectfully disagrees with applicant's interpretation

of the Adya reference. Adya does not apply the restriction to remove indexes from the set of candidate indexes, but instead uses the restriction to specify which indexes can be placed in the candidate index set. In particular, Adya's use the restriction to form a candidate index set in which "all existing indexes are retained" and indexes which are identified by parsing queries in workload are selectively added to the candidate index set. Adya's candidate index set is therefore similar to applicant's candidate index set because it is formed by the union of a current index set and proposed index set identified by parsing queries. In view of applicant's specific example, if Adya's existing indexes set contains {A, B} and the proposed index set contain {C,D}, then the candidate index set contains {A, B,C, D}, same result as seen in applicant's example.

Regarding the 103 rejection based on Chaudhuri (I) and Chaudhuri (II), Applicant argued that Chaudhuri (I) and Chaudhuri (II) do not teach: "forming an index superset from a union of a current index set and a proposed index set". On the contrary, Chaudhuri (I) teaches the storage of "what-if indexes" (i.e., proposed indexes that are absent from the database server) and exiting indexes (i.e., "current index set"). The storage of the indexes serves as an index superset. The examiner agrees with applicant that storing two sets together may be different from union of two sets, because the union set does not contain duplicate entries. However, Chaudhuri (I)'s "what-if indexes" also does not contain duplicate entries because it is formed by exiting indexes and proposed indexes absent from the database server. An index is either existing index or proposed index, but cannot be both. Therefore, Chaudhuri (I)'s "what-if indexes"

comprises two disjoint set, namely current indexes set and proposed index set, and is therefore same as union of current indexes and proposed indexes as claimed.

Applicant argued that "just because storing the element of two disjoint set into a single set yield a result that is the same had the set been combined by taking a union of the sets does not make the storing of the sets the same operation as forming the union of the set", the examiner respectfully submit that if two different operations take the same pair of inputs and produce the same output, then the two operations are equivalent.

In light of the foregoing arguments, the 35 U.S.C 102 and 103 rejections are hereby sustained.

Conclusion

1. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of

the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Khanh B. Pham whose telephone number is (571) 272-4116. The examiner can normally be reached on Monday through Friday 7:30am to 4:00pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Hosain Alam can be reached on (571) 272-3978. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Khanh B. Pham
Examiner
Art Unit 2166

March 23, 2006

